Underwater Sonar Data Fusion Using an Efficient Multiple Hypothesis Algorithm

John J. Leonard
jleonard@mit.edu

Bradley A. Moran
bamoran@mit.edu

MIT Sea Grant College Program
292 Main Street, E38-300
Cambridge, MA 02139
U.S.A.

Ingemar J. Cox
ingemar@research.nj.nec.com

NEC Research Institute
4 Independence Way
Princeton, NJ 08540
U.S.A

Matthew L. Miller
mlm@bimage.tipas.lt.ee

Pilies 32-10
Vilnius
Lithuania

Abstract

This paper describes a geometric approach to underwater environmental modeling using sonar. We classify and localize geometric features of man-made objects by combining the boundary constraints of sonar returns obtained from multiple vantage points. The approach builds on our previous use of Reid's multiple hypothesis tracking (MHT) algorithm in order to resolve data association and motion correspondence ambiguities thereby to construct a model of the observed environment [2]. In particular, we describe a new, computationally efficient implementation of the MHT algorithm originally reported in [3] and validate target models previously developed for air sonar. The technique fuses data by modeling the physics of underwater sonar and its interaction with different object features. We illustrate the approach in two dimensions with real acoustic data taken using a high-frequency (1.25 MHz) pencil-beam profiling sonar, manually positioned along trajectories which circumnavigate prismatic objects.

1 Introduction

Significant interest exists in applying autonomous underwater vehicles (AUVs) to perform useful missions in the harsh ocean environment [20], e.g., to automate the search for hazardous objects lost on the ocean floor. The principal difficulty in automating the object search and recovery process lies with interpreting the data obtained from the vehicle's sensors. While a variety of optical sensors are available for underwater use [16], sonar is a natural choice for investigating the turbid marine environment due to its superior propagation characteristics.

Sonar data can be notoriously difficult to interpret because acoustic returns originate through a variety of physical mechanisms such as specular reflection from a smooth surface and diffuse reflection from a rough surface. Spurious measurements due to multiple reflections are common. Considered in isolation, an individual sonar return yields insufficient information to determine the shape of an object.

Previous underwater research has investigated a variety of sonar interpretation approaches. *Image processing* approaches, for example, aim to imitate the methods by which skilled human operators interpret visual sonar displays [17]. Feature extraction techniques from computer vision, such as edge detection, directly apply in such approaches. Unfortunately, computer vision historically proves extremely difficult to automate. *Signal processing* approaches, in contrast, extract features directly from received signals. One target recognition technique under study, for example, identifies resonant signatures in the echo waveform that correspond to objects of interest [5]. While resonant signatures may aid in the detection of gross material and structural properties, they fail to provide detailed shape information needed for tasks such as grasping. In the latter context, geometric approaches are superior.

Stewart has successfully applied *stochastic backprojection* to a variety of data sets from real ocean settings, including profile data from the USS *Monitor* [15]. This geometric technique employs a grid-based, volumetric representation, similar to the occupancy grid of Elfes [6]. Whereas our approach assumes high accuracy positioning information is available (perhaps through an acoustic transponder system [10]), Stewart explicitly incorporates navigation uncertainty by blurring the measurements over several cells.

IEEE International Conference
on Robotics and Automation
0-7803-1965-6/95 $4.00 ©1995 IEEE

The question of whether a discrete or a continuous geometric representation is more appropriate depends directly on the requirements of the task. For a human-in-the-loop application such as piloting a remotely-operated vehicle (ROV), an occupancy grid offers a method for displaying a large amount of data to the human operator. For object grasping and autonomous recognition of man-made objects, however, an occupancy grid would entail an unduly small cell resolution size, resulting in prohibitive storage and processing requirements. A continuous geometric approach offers the potential of a compact representation capable of efficiently characterizing a scene and providing direct input to higher level reasoning schemes. This can only be possible, however, if the noise and ambiguity of the input data can be handled in a computationally efficient manner.

In this paper, we propose to build accurate geometric reconstructions by grouping together sonar measurements obtained from different vantage points. To accomplish this, we employ the multiple hypothesis tracking (MHT) algorithm, originally developed by Reid [14] in the context of traditional multitarget tracking, and recently shown to be an extremely powerful technique for building and maintaining maps of a robot's environment [2]. We experimentally validate the sensor models previously developed for air sonar [9] with underwater profiling data and describe a new, computationally efficient implementation of the MHT originally reported in [3].

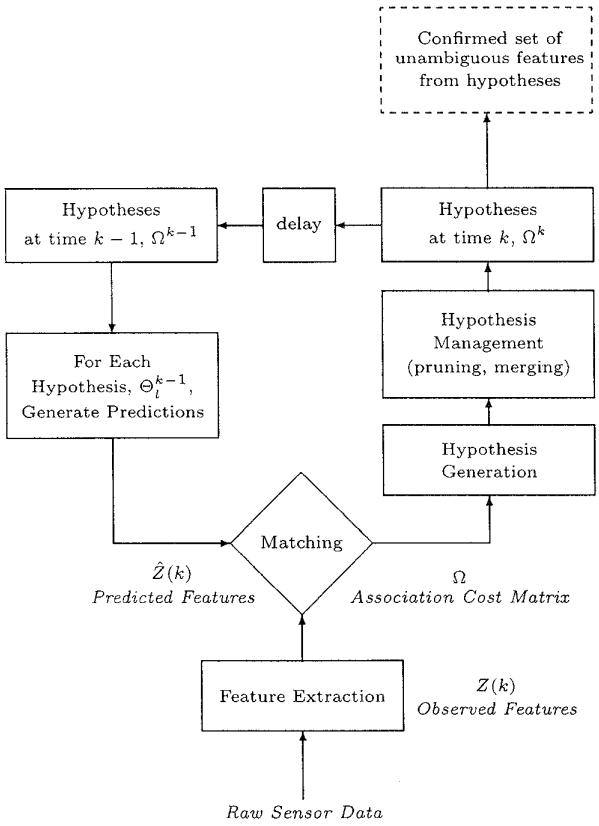

Figure 1: Overview of the MHT algorithm.

2 Review of the MHT Algorithm[1]

The MHT algorithm generalizes state estimation theory to accommodate uncertainty in the origins of measurements. Traditional state estimation techniques such as the Kalman filter apply only to noisy measurements of *known origin*. It provides no resolution to correspondence ambiguities and may even become unstable when given erroneous measurements. The MHT algorithm, in contrast, provides a Bayesian framework for resolving the motion correspondence and data association ambiguities ubiquitous in robotic sensing.

Figure 1 shows the structure of the MHT algorithm [2]. Cycles of the algorithm begin with a list of hypotheses, each denoting a different set of assignments between measurements and features. Each hypothesis thus predicts a set of expected sensor measurements which are compared with actual observed data. Comparisons are represented in the form of an *association cost matrix* which concisely models the ambiguities present in assigning measurements to features. Each hypothesis in the tree has an associated matrix from which it is possible to generate a set of children, each child representing one possible interpretation of the new set of measurements. Containing the growth of the tree necessitates pruning unlikely branches. In order to do this intelligently, we need to evaluate the likelihood of each hypothesis. Reid [14] provides a statistical framework by which to do this under a particular set of assumptions regarding the sensors and the scene.

The MHT is attractive for sensor fusion because it defers making assignment decisions until more data becomes available. Hypotheses containing predictions well supported by subsequent data gain favor relative to hypotheses containing unsupported predictions. By looking ahead two or three time steps and examining the probabilities at the leaves of the tree, the MHT makes very accurate assignment decisions.

Two major difficulties arise in the implementation

[1]This description is a summary of that found in [2].

of the MHT algorithm. It is necessary to formulate sensor and target models for the application domain. In particular, application of the technique in robotics requires multiple target models to handle different types of scene geometry and a physically-based model for the sensor employed. This is discussed in Section 4. Secondly, generating and evaluating so many hypotheses potentially overwhelms computational resources. A solution to this problem is discussed next.

3 MHT Implementation

Several implementation strategies were employed in order to contain the growth of the hypothesis tree and reduce the number of hypotheses that must be considered. In particular, we used the concepts of (1) track trees, (2) spatially disjoint hypothesis trees, and (3) pruning based on an N-scan back procedure, all of which are described by Kurien [8]. However, the most significant implementation strategy was the use of an algorithm due to Murty [12] to generate efficiently the k-best hypotheses directly as orginally described in [3] and briefly discussed next.

Consider first the problem of finding the single most probable hypothesis. This can be cast as a weighted bipartite matching problem by constructing a bipartite graph in which each node on one side represents one of the measurements, each node on the other represents one of the targets, and each arc, $< z_i,\ t_j,\ l >$, gives the log likelihood, l, that measurement z_i should be assigned to target t_j. The log of the likelihood of a given assignment can be found by summing the log likelihoods of all the arcs that it specifies. Finding the best hypothesis, then, is a matter of finding the assignment that maximizes this sum. An instance of the classical assignment problem from combinatorial optimization, it can be solved very efficiently in polynomial time [13].

Murty's [12] algorithm optimally finds the k-best assignments in polynomial time, $O(N^4)$. Furthermore, it finds them in decreasing order of likelihood and is linear in k, a substantial improvement upon previously suggested methods. The application of Murty's k-best assignment algorithm avoids a brute force enumeration of *all* possible global hypotheses, which is entirely impractical, while guaranteeing that the k-best hypotheses are examined. The reader is directed to [3] for more detail.

4 Sensor & Geometric Feature Models

Our profiling sonar mechanically scans through a full 360° panorama and records 400 equally spaced range measurements. It transmits a 1.25 MHz acoustic pulse of 60 μs duration. The circular piston transducer has a radius of 25 mm, with the first null of the corresponding beam pattern at approximately 1.65°. Sufficient energy exists off-axis, however, for some targets to appear through extents up to 20°.

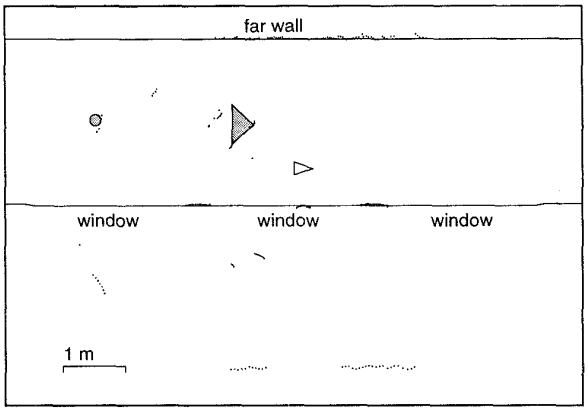

Figure 2: Two prismatic objects, an 18 cm diameter aluminum cylinder and a flooded triangle with thin aluminum walls, stand on end in a 1.3 m deep fresh water tank. The far (concrete) wall of the tank, rough relative to the wavelength of the 1.25 MHz sensor, contrasts with the (glass) windows in the foreground. The unfilled triangle represents the location of the sonar.

Figure 2 shows a typical scan from the sensor taken in the MIT Testing Tank facility. The rear wall of the tank is a diffuse reflector because its surface is rough relative to the 1.2 mm wavelength of the sensor. In contrast, the glass observation windows of the tank are smooth, i.e., specular reflectors. The two objects in the tank are a triangular prism, made by folding aluminum sheets, and an aluminum cylinder with thick walls. Mounted midway in the water column, the sensor scans horizontally such that the scene can be safely approximated as two dimensional.

We have developed a feature extraction technique that detects specular reflections at normal incidence while rejecting many multiple reflections. It assumes an environment of isolated targets, arising from either specular reflection or (knife edge) diffraction; rough surfaces violate this assumption because they give rise to extra returns at high angles of incidence.

4.1 Regions of Constant Depth

Isolated targets, such as a smooth face, a corner, or an edge, produce returns spanning an angular extent determined by the effective beamwidth of the sensor.

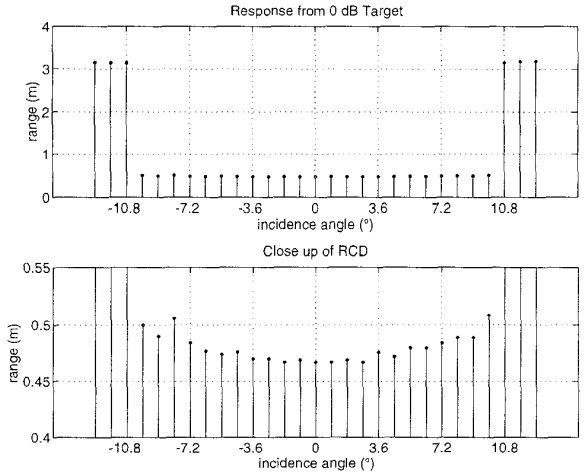

Figure 3: Returns from a target with a smooth face. *Top*, returns appear to have constant range. *Bottom*, in a close up view, however, returns vary in range because the peripheral ones experience delay due to their reduced signal strength.

With an ideal detector, each return in the set would have a range equal to the the distance from the sensor to the *contact point* on the object where a normal from the surface intersects the sensor axis. Figure 3 illustrates this phenomena. The returns associated with the glass window all *appear* to have identical range, hence the origin of the term *region of constant depth* (RCD) to describe this type of feature [9]. A closer look at the measurements using a much finer scale, shown in the bottom half of Figure 3, reveals that the returns are not exactly constant, but rather vary in range by a small amount. The detection delay associated with the lower signal strength found in the peripheral returns causes the feature to deviate from a circular arc as the incidence angle increases. This effect makes the arcs appear "inverted" in Figure 4.

Since individual returns from a target vary in range by a small amount, we introduce a tolerance parameter τ to compare consecutive measurements. An RCD contains a set of consecutive returns which differ in range from their adjacent neighbors by less than this tolerance. A second parameter, w, establishes the minimum angular width an RCD must span. We can choose w small enough to allow only two returns or we can increase it to establish a more stringent extraction filter. The range of an RCD is defined as the minimum range of its member returns, because they correspond to the early detections, implying strong echo (and confidence) levels. The net bearing angle of the RCD is

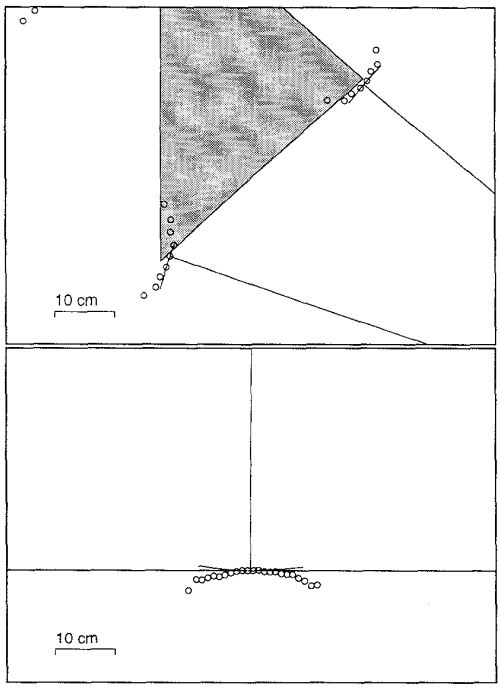

Figure 4: *Top*, sharp edges produce 9 returns each. *Bottom*, a smooth face produces 23 returns. The number of returns generally correlates to the strength of the target. Arcs indicate the angular extent of the RCDs for $\tau = 1$ cm and the rays connecting them to the sensor location indicate the net bearing angle.

the angular centroid of the measurements with minimum range.

Figure 5 shows the resulting RCDs extracted from the scan of Figure 2 using the values of $\tau = 1$ cm and $w = 3.6°$. Note that RCD extraction has greatly reduced the amount of data to consider, but that some spurious RCDs still exist. One of these RCDs seems to originate from behind the triangle. In fact this RCD results from transmission through the near face of the triangle, specular reflection off the remaining two faces, and finally transmission back through the front face to be detected at the sensor. Unlike with an air sonar such as the Polaroid sensor, in water one sees a much wider variation in the acoustic reflection coefficient of objects; partial object transparency is common and is an additional source of spurious returns that must be rejected by the data association algorithm.

The feature detector finds two RCD's on the diffuse wall. One is due to the set of returns near normal incidence, and presents the same geometric constraint as for the smooth glass window. The other is due to

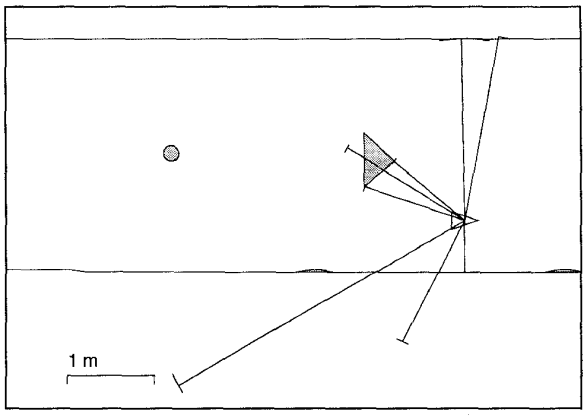

Figure 5: Sample RCDs extracted using typical parameters ($\tau = 1$ cm and $w = 3.6°$).

returns at a higher angle of incidence, the consequence of which is discussed within the next section.

4.2 Geometric feature models

For polygonal and cylindrical scenes, the geometric features of objects can be represented as points in a parameter space and estimated with a Kalman filter. As described in [2], at time k, a 2D line is represented by the state vector $\mathbf{x}_t(k) = [R_t(k), \theta_t(k)]^T$, where $R_t(k)$ is the perpendicular distance from the origin of the global coordinate frame to the line and $\theta_t(k)$ is the orientation of a perpendicular drawn from the line to the origin. In addition, each line has two endpoints associated with it, which are estimated by projecting sensor locations onto the infinite line defined by $[R_t(k), \theta_t(k)]^T$. A corner is represented as a point in 2D by $\mathbf{x}_t(k) = [x_t(k), y_t(k)]^T$. The location of the vehicle is specified by the vector $\mathbf{x}_R(k) = [x(k), y(k), \theta(k)]^T$. The polar measurements $\mathbf{z}_j(k) = [r_j, \alpha_j]^T$ are a function of $\mathbf{x}_t(k)$ and $\mathbf{x}_R(k)$, subject to a noise disturbance $\mathbf{w}_t(k)$, as given by the measurement model

$$\mathbf{z}_j(k) = \mathbf{h}_t(\mathbf{x}_R(k), \mathbf{x}_t(k)) + \mathbf{w}_j(k)$$

where the measurement function $\mathbf{h}_t(\cdot, \cdot)$ takes a different form depending on the type of geometric feature. For a plane,

$$r = |R_t - x(k)cos(\theta_t) - y(k)sin(\theta_t)|$$

$$\alpha = \theta_t(k) - \theta(k),^2$$

while for a corner,

$$r = \sqrt{(x_t - x(k))^2 + (y_t - y(k))^2}$$

$$tan(\alpha + \theta(k)) = \frac{y_t - y(k)}{x_t - x(k)}, \qquad x_t \neq x(k).^3$$

The measurement noise is modeled by zero-mean Gaussian noise with standard deviation $\sigma = 1$ cm in range and $\sigma = 4$ degrees in angle. The multiple hypothesis algorithm, described next, uses these models to compute recursively the locations of geometric features, based on hypothesized assignments of the data.

5 Underwater Shape Reconstruction

Having isolated RCDs (circular arcs) from the raw sonar data, we need to determine which arcs originate from which geometric features/objects, and which arcs to consider spurious and ignore. We accomplish this by moving the sensor and *tracking* the RCD features from one location to the next. Each new sonar measurement provides an additional geometric constraint on the class of possible 2D surfaces being viewed. Since a specular return implies a surface tangent to the RCD, multiple tangency constraints from adjacent sensing locations provide a sequence of "contact points" to trace out an object's surface. Returns from a point target such as an edge or a corner provide the geometric constraint that the target lies at the intersection of the RCDs.

The above formulation is a multiple hypothesis tracking problem. At each iteration or scan, measurements in the form of extracted RCDs are obtained. The MHT compares these measurements with predictions based on the possible target models: (1) a point feature such as a corner, (2) a specular planar surface, and (3) a specular cylindrical surface [11]. Each measurement may (1) belong to an existing geometric feature, (2) initiate a new geometric feature (of which exist three types), or (3) be a spurious measurement such as a multiple reflection. In addition, existing geometric features (tracks) that fail to be supported by a measurement may either be continued to the next scan or terminated, i.e., they have exited the field of view of the sensor.

Figure 6 shows the RCD features extracted, using $\tau = 1$ cm and $w = 3.6°$, during a complete sequence around the triangular object. The sensing locations were hand-measured to within a few millimeters of accuracy. A typical vertex track for this object, shown along with the location estimate in Figure 7, contains 16 RCDs. Figure 8 shows all six of the geometric features superimposed on the triangular object to which they belong. The estimated geometric features for the faces of the triangle are shown as infinite lines.

[2]More precisely, we must occasionally add π to the angle in order to resolve on which side of the plane lies the sensor.

[3]In practice, IEEE arithmetic supports the degenerate case when $x_t = x(k)$.

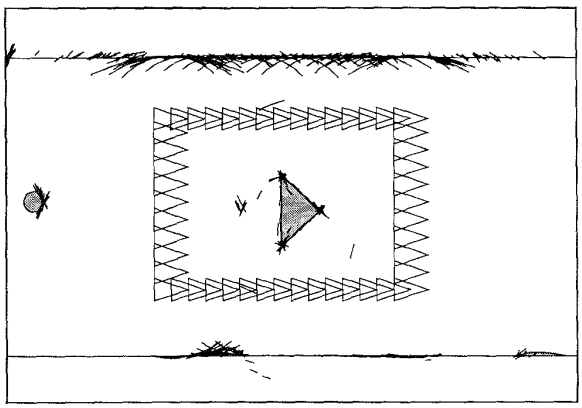

Figure 6: RCD features extracted from a complete sequence around a triangular object.

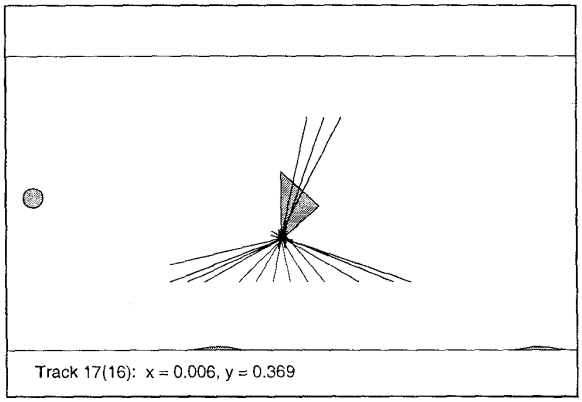

Figure 7: A typical vertex track, which comprises 16 RCDs determined to originate from the same vertex and an estimate of its location.

Two of these tracks contain three RCDs and the third track (lower right face) contains four. The three vertex tracks (clockwise from the top) each contain 12, 13, and 16 RCDs. The specular alignment condition needed to observe the faces of the triangle causes fewer extracted RCDs. Vertex estimates lie inside the true locations because of the delay in the threshold detector of our sensor. Targets that are acoustically stronger (because of shape and/or material properties) will appear closer than weaker targets [19]. Table 1 summarizes the feature locations and their corresponding errors. Although this discussion focuses on the triangular object in Figure 2, we have applied the algorithm to a variety of 2D scenes in the MIT Testing Tank [11].

Figure 9 shows a superposition of the remaining tracks. We see many point targets that lie "behind" the top rough wall. Because the wall is not an isolated

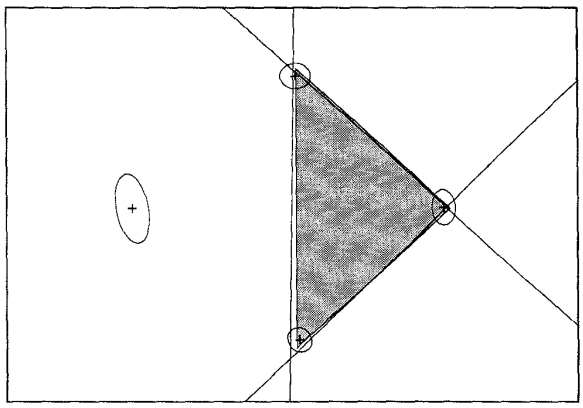

Figure 8: Close-up view of feature estimates for the triangular object. 4σ error ellipses are show for vertex features. The feature to the left is a reflection of the interior corner of the right vertex in the left face.

Track	Faces					
	Estimated		Actual		Difference	
	θ (deg)	$R(m)$	θ (deg)	R (m)	θ (deg)	R (m)
left	1.02	-0.009	0	0	-1.02	0.009
lower	-133.61	-0.283	-132.17	-0.286	1.44	-0.003
upper	132.13	-0.166	132.36	-0.171	0.23	-0.005
	Vertices					
	x (m)	y (m)	x (m)	y (m)	x (m)	y (m)
lower	0.006	0.369	0	0.386	-0.006	0.017
upper	-0.001	-0.215	0	-0.232	0.001	-0.017
right	0.326	0.076	0.34	0.078	0.014	0.002

Table 1: Comparison of estimated and hand-measured feature locations for the triangular object. Faces are parameterized by their distance R from the origin and their orientation θ.

target it produces extra RCDs at high angles of incidence. A good track is still produced for the rough wall from the normal-incidence RCDs. Figure 10 shows the measurements rejected as false alarms.

The cumulative data set for this experiment contained 48 scans with an average of approximately 6 RCDs per scan. The false alarm likelihood was set to 0.1 and the detection likelihood was set to 0.45. On A 50 MHz RISC workstation, the MHT algorithm required approximately 5 seconds to process all of the data using the following values for pruning parameters: N = 4-scan-back, k = 500-best, and minimum likelihood ratio 0.01. See [2] and [11] for further discussion of the role of these parameters.

6 Discussion

Quite accurate reconstructions are possible because the algorithm groups together returns obtained from different vantage points that originate from the same

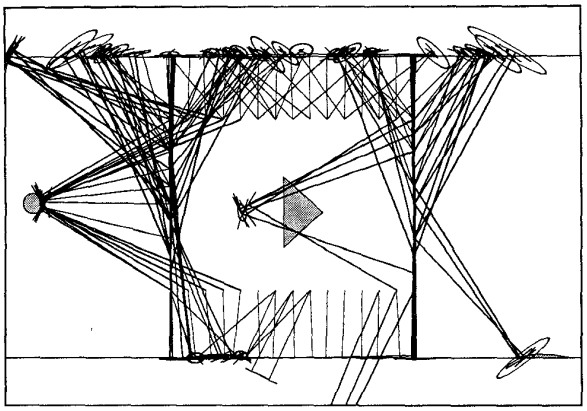

Figure 9: Remaining tracks not shown in Figure 8, displayed with RCDs. An erroneous vertex track caused by multiple reflections is prominent in the middle of the figure. This track is the "image" of the right vertex of the triangle reflected by the left face of the triangle. There are a large number of point features initialized at high angles of incidence to the rough top wall of the tank, because it violates the assumptions of the RCD extractor; these all fall behind a single long plane track created for the wall from its normal-incidence responses. The arcs from the cylinder are grouped into a single point target.

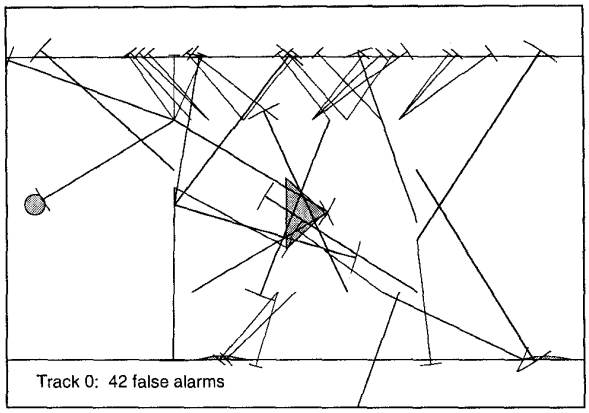

Figure 10: Features classified as false alarms.

geometric feature, while rejecting spurious measurements. Our approach addresses uncertainty in the *origins* of measurements (data association or correspondence uncertainty) as well as in the *values* of measurements (noise uncertainty). It models the uncertainty due to noise by covariance matrices and represents the data association uncertainty by Bayesian probabilities attached to nodes of a hypothesis tree, each node representing different possible assignments of measurements to features. In addition to sonar interpretation, the MHT approach has been applied to a surprising number of diverse areas such as motion tracking [1], contour grouping [4], and seismic data interpretation [18].

The experimental results for the triangle scene demonstrate that the geometric sonar interpretation approach developed for land mobile robotics by Kuc [7], Leonard [9] and others can be applied in an underwater setting. Further, the results validate the efficient hypothesis generation strategy described in Section 3.

Because the experiments thus far have been restricted to static, rigid, 2D scenes, further research is necessary to fulfill the task of underwater object search and retrieval we described at the outset. Theoretically, the techniques readily extend to 3D (arc features, for example, transform into solid angle sections of a sphere). In practice, *track initiation* certainly becomes more problematic.

The MHT is formulated based on the assumption that the target state estimate (in our case the location of a face, edge, or vertex) can be initialized from a single measurement. Subsequent measurements are evaluated via data-to-target association hypotheses, by comparing new observations to predictions made based on the estimated target state. In three dimensions, however, target states contain too many degrees of freedom to be initialized from a single measurement. Two solutions appear feasible. The first is to remove responsibility for track initiation from the MHT and provide a specialized module for the purpose, much as Kurien [8] has done. The second solution is employ an array of sonar sensors to make several measurements of a geometric feature at each scan, thereby providing sufficient information to allow the initiation of geometric features.

Finally, the specular wavelength regime only applies to objects whose RMS surface roughness is less than $\lambda/4$. Under typical marine conditions of both sedimentation and biological fouling, many surfaces will reflect diffusely, not specularly. As such, they will cease to fit the conditions of being an isolated target. With a feature extraction technique that can adapt to a range of surface roughness scales, it should be possible to classify a diffuse scatterer such as the top wall in Figure 2 as a single track.

Acknowledgements

We would like to thank Prof. C. Chryssostomidis and Dr. J. G. Bellingham, both of MIT. This work

was partially supported by the Advanced Research Projects Agency under contract MDA 972-88-C-0040.

References

[1] I. J. Cox and S. L. Hingorani. An efficient implementation and evaluation of reid's multiple hypothesis tracking algorithm for visual tracking. In *Int. Conf. on Pattern Recognition*, pages 437–442, 1994.

[2] I. J. Cox and J. J. Leonard. Modeling a dynamic environment using a Bayesian multiple hypothesis approach. *Artificial Intelligence*, 66:311–344, 1994.

[3] I. J. Cox and M. L. Miller. On finding ranked assignments with application to multi-target tracking and motion correspondence. *IEEE Trans. on Aerospace and Electronic Systems*, 32(1), 1995.

[4] I. J. Cox, J. M. Rehg, and S. Hingorani. A Bayesian multiple hypothesis approach to contour grouping and segmentation. *Int. J. of Computer Vision*, 11(1):5–24, 1993.

[5] C. Dean and M. Werby. Target shape and material composition from resonance echoes of submerged elongated elastic targets. In *Proc. SPIE Conf. on Automatic Object Recognition*, volume 1471, pages 54–65, Orlando, FL, 1991.

[6] A. Elfes. Sonar-based real-world mapping and navigation. *IEEE J. of Robotics and Automation*, RA-3(3):249–265, June 1987.

[7] R. Kuc and M. W. Siegel. Physically based simulation model for acoustic sensor robot navigation. *IEEE Trans. Pattern Analysis and Machine Intelligence*, PAMI-9(6), November 1987.

[8] T. Kurien. Issues in the design of practical multitarget tracking algorithms. In Y. Bar-Shalom, editor, *Multitarget-Multisensor Tracking: Advanced Applications*, pages 43–83. Artech House, 1990.

[9] J. J. Leonard and H. F. Durrant-Whyte. *Directed Sonar Sensing for Mobile Robot Navigation.* Kluwer Academic Press, Boston, 1992.

[10] Marquest Telepresence, Bourne, MA. Commercial offer: EXACT acoustic navigation system, 1992.

[11] B. A. Moran. *Underwater Shape Reconstruction in Two Dimensions.* PhD thesis, Massachusetts Institute of Technology, May 1994. Supervised by C. Chryssostomidis.

[12] K. G. Murty. An algorithm for ranking all the assignments in order of increasing cost. *Operations Research*, 16:682–687, 1968.

[13] C. H. Papadimitriou and K. Steiglitz. *Combinatorial Optimization.* Prentice-Hall, 1982.

[14] D. B. Reid. An algorithm for tracking multiple targets. *IEEE Trans. on Automatic Control*, AC-24(6):843–854, December 1979.

[15] W. K. Stewart, Jr. *Multisensor Modeling Underwater with Uncertain Information.* PhD thesis, Massachusetts Institute of Technology, 1988.

[16] W. K. Stewart, Jr. Remote sensing issues for intelligent underwater systems. In *Intl. Conf. on Computer Vision*, pages 230–235, 1991.

[17] G. Trimble, J. Vilaro, D. Okamura, R. Lum, and K. Dutta. Underwater object recognition and automatic positioning to support dynamic positioning. In *Proc. of the 7th Intl. Symposium on Unmanned Untethered Submersible Technology*, pages 273–279, 1991.

[18] P. Tu, A. Zisserman, I. Mason, and I. J. Cox. Identification of events from 3D volumes of seismic data. In *Int. Conf. on Pattern Recognition*, pages 128–130, 1994.

[19] R. J. Urick. *Principles of Underwater Sound.* McGraw-Hill, 3d edition, 1983.

[20] G. Zorpette. Autopilots of the deep. *IEEE Spectrum*, 31(8):38–44, August 1994.